Opportunities in scaleable, high-performance interconnects

- Desire to interconnect of order 10^4 CPUs, or 10^3 computing nodes
- Balanced 10-to-30 TFLOPS system in late 1999-to-2001 requires 5-to-15 TBytes/sec bisection bandwidth
- Desired latency is a few hundred CPU cycles

Examples of areas of interest

- Scaleable interconnect prototype development and demonstration
  - not necessarily at full scale
  - may include node-to-interconnect adapter
- Standardized, scaleable node-to-interconnect adapter interface
- Advanced interconnect features
  - extension of cache-coherence domain
  - large data set migration; cache-only memory architecture

ASCI hardware requirements

Level	Effective Latency (CPU cycles)	Bandwidth (Random read/write)	Size	
On-chip cache**, L1	2-3 ●	16-32 B/cycle ●	10^{-4} B/flop * ● ↑	
Off-chip cache**, L2 (SRAM)	5-6 ●	16 B/cycle ●	10^{-2} B/flop * ● ↑	
Local main memory (DRAM)	30-80 (15-30) ↓	2-8 B/flop pk (2-8 B/flop sustained) ↓	1 B/flop ● ↑	Compute engine
"nearby nodes"	300-500 (30-50) ↓	1-8 B/flop (8 B/flop) ↓	1 B/flop ●	Interconnect
"far away nodes"	1000 (100-200) ↓	1 B/flop (1 B/flop) ↓	1 B/flop ●	
I/O (memory disk)	10 ms ●	0.01-0.1 B/flop ●	10-100 B/flop ●	
Archive (disk-tape)	Seconds ●	10^{-4} B/flop (0.001-0.01 B/flop) ↓	10^{2} B/flop 10^{4} B/flop ↓	
User access	1/10 s (1/60 s)	OC3/desktop (OC12-48 /desktop) ↓	100 users ●	
Multiple sites	1/10 s ●	●	●	

* Equivalent integer and floating-point data calculation rates are required.
** Cacheless systems with equivalent performance are fully acceptable.

Primary investment priority

Secondary investment priority

1996-1998 Situation

(1998-2000 Requirements)

Industry Trend

↑ Industry gets better at meeting requirements

↓ Industry gets worse at meeting requirements

● Industry continues to meet requirements